Flux, Surface Integrals & Gauss' Law

A Guide for the Perplexed

0. Introduction

What I want to do tonight is

- Define the concept of "flux", physically and mathematically
- See why an integral is sometimes needed to calculate flux
- See why in 8.02, you'll almost never need an integral to calculate flux ☺
- Go through some examples
- See how it relates to Gauss' Law, and go through some more examples

1. What is flux?

The aim of a **surface integral** is to find the **flux** of a **vector field** through a **surface**. It helps, therefore, to begin what asking "what is flux"? Consider the following question

> "*Consider a region of space in which there is a constant vector field,* $\boldsymbol{E}(x,y,z) = a\hat{\boldsymbol{x}}$ *. What is the flux of that vector field through an imaginary square of side length L lying in the y-z plane?*"

As ever, let's begin with a diagram:

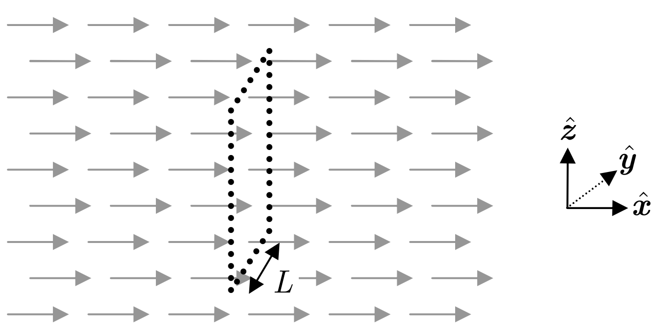

So, how are we going to answer the question? It turns out that a very useful way to think of the flux is as follows:

- Imagine that the vector field is a fluid flowing through space

- Then, the **flux** of the field through an area is the **amount of "fluid" flowing <u>through</u> that area**.

Now, in this case, the area we're flowing through is L^2, and the field strength is a, and so

$$\boxed{\text{Flux} = aL^2}$$

Now, consider a slightly different example

> "*Consider a region of space in which there is a constant vector field,* $\boldsymbol{E}(x,y,z) = a\hat{\boldsymbol{x}}$. *What is the flux of that vector field through an imaginary square of side length L lying in the x-y plane?*"

Once again, a diagram

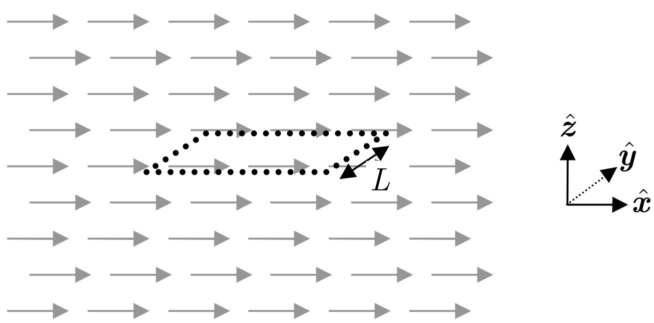

Let's think back to our fluid analogy. In this case, it should be pretty clear that *none* of the fluid is passing through our square, because the plane of the square is **parallel** to the fluid flow. It's like trying to blow a bubble with the bubble hoop turned the wrong way – you won't get any air flowing through the hoop, and so no bubble ☹. As such, in this case

$$\boxed{\text{Flux} = 0}$$

2. The vector $\hat{\boldsymbol{n}}$

The examples we've looked at so far are pretty easy – almost trivial – because the square was either perpendicular or parallel to the flow. In more complicated examples, the area could be along any arbitrary direction. To be able to cope with such cases, we need to develop a mathematical way to describe the orientation of an area.

The tool we're going to use to do this is called the **normal vector** (denoted $\hat{\boldsymbol{n}}$) to the area. It is a **unit vector** that points in a direction **perpendicular to the area in question**. So, for example, for the two planes we considered above, the normal vectors are:

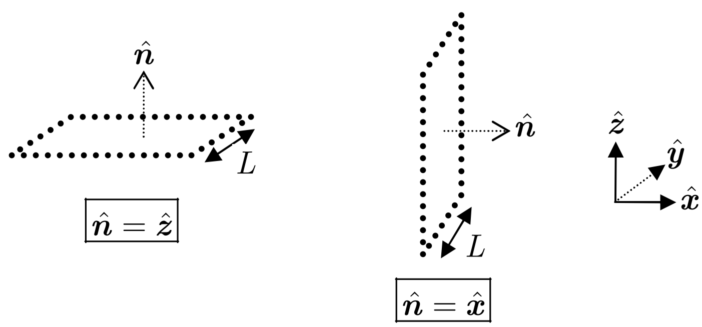

The more eagle-eyed amongst you will have noticed that there is some ambiguity as to how $\hat{\boldsymbol{n}}$ is defined. In the first case, I could have chosen it downwards instead of upwards, and in the left case I could have chosen it to the left instead of to the right. This is a problem we'll talk more about later.

Note: Some texts choose to define a vector $\boldsymbol{A} = \text{Area} \times \hat{\boldsymbol{n}}$. For pedagogical reasons, I'm not going to do this in this handout. But I don't want it to throw you off if you come across it in a book…

In any case – to return to the main theme of this evening's symposium – we're not equipped with a way to deal with much more bizarre planes…

3. Planes in arbitrary directions

…like in this problem

> "*Consider a region of space in which there is a constant vector field of magnitude a, in a direction inclined at an angle θ to the normal vector of a square of side length L. What is the flux of the field through the square?*"

Again, let's get to ball rolling with a diagram:

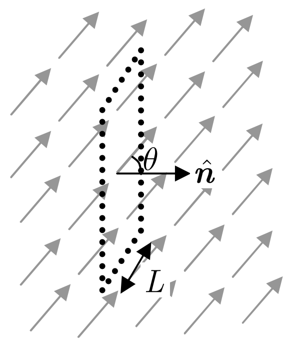

How can we proceed with this problem? The key realisation comes when we notice that the field vectors can be broken down into two components, as follows (for clarity, I'll only draw one vector):

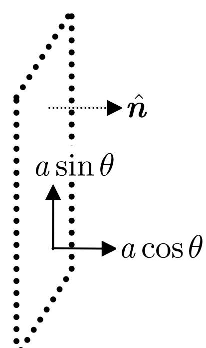

Now, let's have a look at the flux of each of these components through the square:

- **The** $a\sin\theta$ **component** is *parallel* to the plane, and therefore, as in the second example above, it will produce **0 flux**.
- **The** $a\cos\theta$ **component** is *perpendicular* to the plane, and therefore, as in the first example above, it will produce a flux $aL^2\cos\theta$

Adding these two fluxes together, we find that the answer to our question is

$$\boxed{\text{Flux} = aL^2 \cos\theta}$$

4. Writing these results using elegant math

You might be starting to see a pattern emerge from the three results above… Consider:

- In the first example, the field was $\boldsymbol{E} = a\hat{\boldsymbol{x}}$ and the normal vector was $\hat{\boldsymbol{x}}$. The total flux was aL^2.
- In the second example, the field was also $\boldsymbol{E} = a\hat{\boldsymbol{x}}$, but the normal vector was $\hat{\boldsymbol{y}}$. There, the total flux was 0.
- In the third example, the field and normal vector had an angle θ between then, and the $\boldsymbol{E}$ vector had magnitude a. The total flux was $aL^2 \cos\theta$

It looks, in fact, like the general expression for flux is

$$\text{Flux} = \left(\boldsymbol{E}\cdot\hat{\boldsymbol{n}}\right)\times\text{Area}$$
$$\boxed{\Phi = \boldsymbol{E}\cdot\hat{\boldsymbol{n}}A}$$

If you think about it, this result makes sense – the dot product basically takes the **component of the vector field** that is **parallel to $\hat{\boldsymbol{n}}$** and so **perpendicular to the plane** and **discards the other component**. Finally, we multiply by the area.

5. Surface "integrals"?

So far so good, but we haven't seen any integrals! What's going on? What's happening is that so far, I've chosen examples with particularly nice properties – let's look at our expression for flux again

$$\boxed{\text{Flux} = \boldsymbol{E}\cdot\hat{\boldsymbol{n}}A}$$

What's so nice about the two examples above is that both $\boldsymbol{E}$ and $\hat{\boldsymbol{n}}$ were constant over the area we were considering. How so?

- $\boldsymbol{E}$ wasn't changing (it was a constant vector field)
- We were considering a simple plane, and so one particular $\hat{\boldsymbol{n}}$ vector was perpendicular to every point in the area we were considering.

Some cases, however, aren't so nice, and both the bits above could be changing:

- $\boldsymbol{E}$ could be changing, if our vector field wasn't constant.

- $\hat{\boldsymbol{n}}$ could be changing, if the surface we're considering isn't flat. For example – consider a hemisphere. Clearly, there is no one $\hat{\boldsymbol{n}}$ vector that's perpendicular to every point on the surface. The direction of the surface charges as we go round

In such cases, we have to do something which you've hopefully seen many times by now:

- We split up the surface into lots of tiny bits
- We find the values of $\boldsymbol{E}$ and $\hat{\boldsymbol{n}}$ for these bits
- We find the area of each bit ($\mathrm{d}A$)
- We the simply perform the integral to find the flux

$$\boxed{\Phi = \int_{\substack{\text{Every bit}\\ \text{in our surface}}} \boldsymbol{E}\cdot\hat{\boldsymbol{n}}\,\mathrm{d}A}$$

Which is the expression for flux that you are probably familiar with. The one last step is a bit of notation – when the surface is closed, we often put a ring around the integral:

$$\boxed{\Phi = \oint_{\substack{\text{Every bit}\\ \text{in our surface}}} \boldsymbol{E}\cdot\hat{\boldsymbol{n}}\;\mathrm{d}A}$$

Our steps in calculating it are always going to be

1. Decide how you're going to split up your surface (see the next section)
2. Check if the flux through any bit of your surface is obviously 0.
3. Find an expression for $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$
4. Check that $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ isn't constant (see later!)
5. Find $\mathrm{d}A$ for each little bit your surface (see next section)
6. Set the correct limits on your integral
7. Solve it!

6. Common ways to split up surfaces

We'll look at lots of examples with fields soon, but before we do, I want to examine this process in step 1 and 5 in more detail – that of "splitting up surfaces". You'll be glad to know that there are really only three possible surfaces you will need to use in 8.02…

1. **Planes stuck together**

   The first, and simplest case, is a case in which lots of planes are stuck to each other – for example, a cube. In such a case, the "bits" you split your shape into are simply the faces of the cube, each with their own $\hat{\boldsymbol{n}}$ vector:

   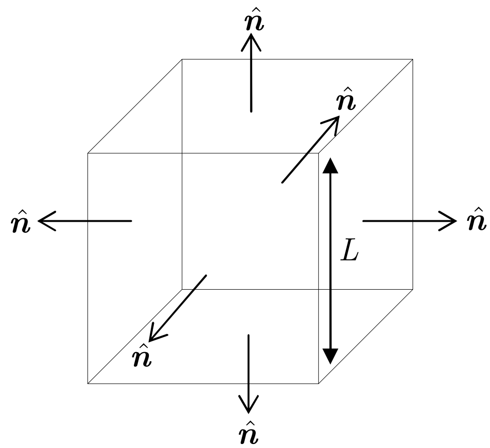

   The area of each face is then

   $$\boxed{\text{Area} = L^2}$$

   We then simply find the flux through each face, and add them all up.

2. **A sphere**

   A sphere is slightly more complicated. What we do in this case is subdivide the sphere into lots of tiles, each of angular thickness $\mathrm{d}\theta$ and $\mathrm{d}\phi$:

   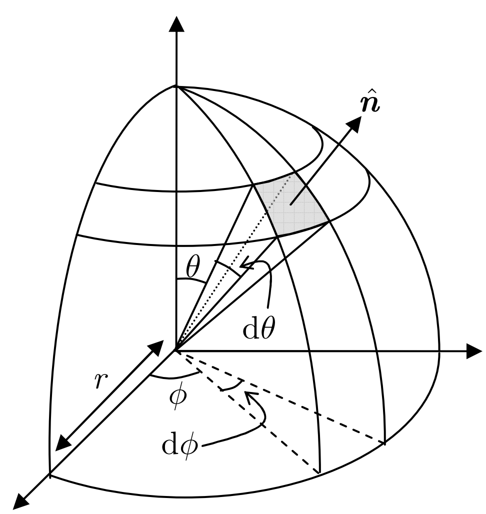

The area of each tile is given by

$$\boxed{\mathrm{d}A = r^2 \sin\theta \,\mathrm{d}\theta \,\mathrm{d}\phi}$$

And in this case, the normal vector from each tile points directly outwards, in the $\hat{\boldsymbol{r}}$ direction, so

$$\boxed{\hat{\boldsymbol{n}} = \hat{\boldsymbol{r}}}$$

The plan to find the total flux is then to

- Find the flux through each little tile, by calculating $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}\,\mathrm{d}A$ for the tile.
- Sum up (= integrate) all these fluxes over the entire surface

3. **A cylinder**

A cylinder is a combination of the two examples above. The top and bottom of the cylinder have simple normal vectors, because they are planes, but the curved sides have to be divided into lots of little tiles, each with their own normal vector:

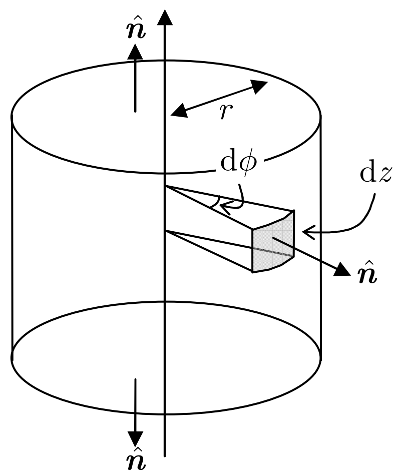

The area of each tile is given by

$$\boxed{\mathrm{d}A = r\,\mathrm{d}\phi\mathrm{d}z}$$

And the area of the top and bottom faces is given by

$$\boxed{\text{Area} = \pi r^2}$$

And once again, the normal vector from each of the tiles clearly points in the $\hat{\boldsymbol{r}}$ direction, so

$$\boxed{\hat{\boldsymbol{n}} = \hat{\boldsymbol{r}}}$$

To find the total flux, we simply find the flux over the curved surface area by integrating, then find the flux through the two ends, and add them all up.

One last point before we move on – you'll notice that I managed to resolve the aforementioned ambiguity in choosing $\hat{\boldsymbol{n}}$ using the following very useful convention:

The $\hat{\boldsymbol{n}}$ vector for any closed surface always point
outwards from the surface

This then defines all the $\hat{\boldsymbol{n}}$ vectors uniquely.

7. Why the integral is often very easy…

So far so good – but you'll be delighted to know that there are a number of factors that often make the integral much, much easier. Our aim in choosing a surface in Gauss' law will be to choose a surface that fulfils one of these conditions

The flux through some of the faces might be 0

This could happen for two reasons:

- $\boldsymbol{E}$ could be perpendicular to $\hat{\boldsymbol{n}}$ on that surface, which means $\boldsymbol{E}\cdot\hat{\boldsymbol{n}} = 0$. This such means that the field is parallel to the surface; the second example we did today.
- $\boldsymbol{E}$ could be 0 on that surface…

$\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ might be constant

This is very handy. Let's look at the integral again

$$\Phi = \int_{\substack{\text{Every bit}\\ \text{in our surface}}} \boldsymbol{E}\cdot\hat{\boldsymbol{n}}\,\mathrm{d}A$$

If you calculate $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$, and find that it's a constant, you can take it out of the integral and you're left with

$$\Phi = \boldsymbol{E}\cdot\hat{\boldsymbol{n}}\int_{\substack{\text{Every bit}\\ \text{in our surface}}} \mathrm{d}A$$

But the integral is now trivial! All it's saying is "add up all the areas of the little bits that make up my surface". The result, of course, will just be the area of the surface. Therefore

$$\boxed{\Phi = \left(\boldsymbol{E}\cdot\hat{\boldsymbol{n}}\right)\times A}$$

This formula might look familiar – it's the very first one we derived above. It was true there because $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ was indeed constant ($\boldsymbol{E}$ was a constant vector field and $\hat{\boldsymbol{n}}$ was the same over the whole plane).

8. Examples, examples, examples…

OK! Time for lots of examples!

> "*Consider a region of space in which there is a constant vector field,* $\mathbf{E}(x,y,z) = a\hat{\boldsymbol{x}}$. *What is the flux of that vector field through a cube of side length L lying perpendicular to the field?*"

As ever, we begin with a diagram

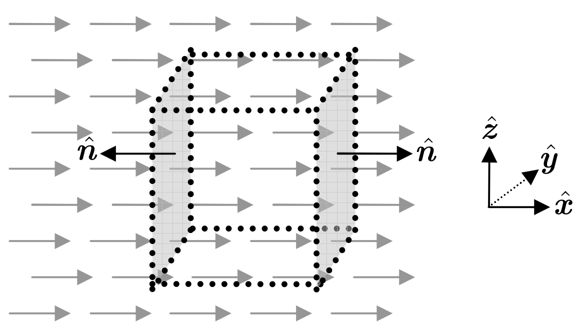

Let's go through the seven steps!

1. **Decide how you're going to split up your surface**

   Clearly, we need to split it into each of the faces, so we have 6 planes

2. **Check if the flux through any bit of your surface is obviously 0**

   It's also clear that only the grey faces above will have flux through them, because the other faces are parallel to the field lines.

3. **Find an expression for $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$**

For both faces, $\boldsymbol{E}(x,y,z) = a\hat{\boldsymbol{x}}$. For the left-hand-face, $\hat{\boldsymbol{n}} = -\hat{\boldsymbol{x}}$, and for the right-hand-face, $\hat{\boldsymbol{n}} = +\hat{\boldsymbol{x}}$. Therefore, for the left-hand-face, $\boldsymbol{E}\cdot\hat{\boldsymbol{n}} = -a$ and for the right-hand-face, $\boldsymbol{E}\cdot\hat{\boldsymbol{n}} = a$ **.**

At his point, we can just stop. The fluxes for those two faces cancel, because $a - a = 0$. So

$$\boxed{\text{Flux} = 0}$$

In retrospect, this makes sense… Every field line that enters the box also leaves it, so no "net flux" is coming into the box.

> "*Consider a region of space in which there is a vector field $\boldsymbol{E}(x,y,z) = +a\hat{\boldsymbol{z}}$ above the x-y plane, and a field $\boldsymbol{E}(x,y,z) = -a\hat{\boldsymbol{z}}$ below the x-y plane. What is the flux of that vector field through a cube of side length L with its centre at the origin?*"

As usual, the diagram (I've tried to draw as few field lines as possible, to make he diagram less cluttered – it shows that in the top half of the cube, the field lines point up, whereas in the bottom half, they point down.

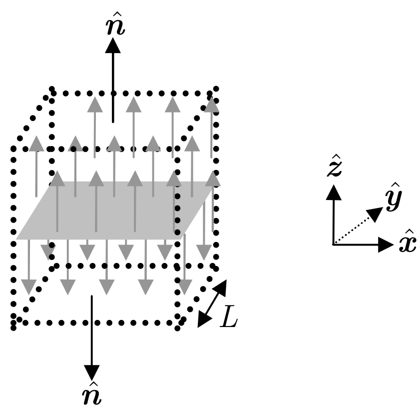

Once again, let's go for our seven steps

1. **Decide how you're going to split up your surface**

Again, we'll split it up into the six faces.

2. **Check if the flux through any bit of your surface is obviously 0**

   Again, flux will only be passing through the top and bottom surface. The side faces are parallel to the field, and so the flux through them is 0.

3. **Find an expression for $\boldsymbol{E} \cdot \hat{\boldsymbol{n}}$**

   For the top face, $\boldsymbol{E} = a\hat{\boldsymbol{z}}$ and $\hat{\boldsymbol{n}} = +\hat{\boldsymbol{z}}$, and so $\boldsymbol{E} \cdot \hat{\boldsymbol{n}} = a$. For the bottom face, $\boldsymbol{E} = -a\hat{\boldsymbol{z}}$ and $\hat{\boldsymbol{n}} = -\hat{\boldsymbol{z}}$, so we also have $\boldsymbol{E} \cdot \hat{\boldsymbol{n}} = a$. This time, the fields do not cancel, because the field changes direction halfway through.

4. **Check that $\boldsymbol{E} \cdot \hat{\boldsymbol{n}}$ isn't constant**

   It's pretty clear that $\boldsymbol{E} \cdot \hat{\boldsymbol{n}}$ is indeed constant over these surfaces. So no integration! Yay!

5. **Find d$\boldsymbol{A}$ for each little bit your surface**

   Each of the faces has area L^2

And that's it! Since $\boldsymbol{E} \cdot \hat{\boldsymbol{n}}$ is constant, we just multiply it by A for each face, and get

$$\text{Flux} = aL^2 + aL^2$$

$$\boxed{\text{Flux} = 2aL^2}$$

Once again, this kind of makes sense – this time, we have field lines being "created" inside the cube, so there's a net flux through its faces.

> "*A planet of radius R rests in a constant vector field $\boldsymbol{E}(x,y,z) = a\hat{\boldsymbol{x}}$. What is the flux through the surface of the planet*"

First, a diagram

This is another example of a situation in which a field goes through a closed surface. Clearly, every field line that enters the planet also leaves it, and so the total flux through the planet is 0.

If you want to think about it mathematically, you can realise that the field lines are always pointing towards the left, but that the normal vectors are always pointing outwards from the planet, and therefore sometimes point left and sometimes point right.

This fact (that fields through closed surfaces give 0 flux) will be extremely useful in our discussion of Gauss' Law. However, a slight caveat – this only applies if the **divergence** of the field is 0 ($\nabla \cdot \boldsymbol{E} = 0$) in the region of the closed surface. If you don't know what that means, don't worry – you'll find out in a future math session.

> "*Consider a cylinder of radius* r *and length* L*. A field* $\boldsymbol{E}(r,\phi,z) = a\hat{\boldsymbol{r}}$ *radiates from the line through the centre of the cylinder. Find the flux of the field through the surface of the cylinder.*

Let's draw a diagram (as usual, the field lines are in grey):

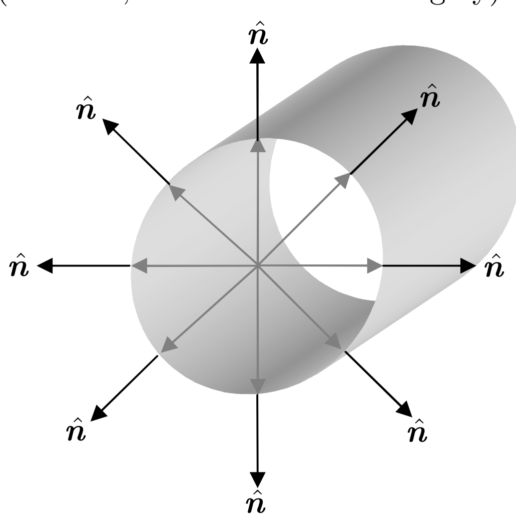

Once again, let's go through the steps:

1. **Decide how you're going to split up your surface**

   In this case, it's pretty clear we have cylindrical symmetry. So we'll say the two ends of the cylinder are planes, and we'll split the curved surface into tiles.

2. **Check if the flux through any bit of your surface is obviously 0**

   The flux through the two ends is obviously 0, because the field is parallel to those ends.

3. **Find an expression for $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$**

   In this case, we saw that $\hat{\boldsymbol{n}} = \hat{\boldsymbol{r}}$, and we told that $\boldsymbol{E} = a\hat{\boldsymbol{r}}$. Therefore, $\boldsymbol{E}\cdot\hat{\boldsymbol{n}} = a$ (this makes sense – the field lines and normal vectors are parallel to each other).

4. **Check that $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ isn't constant**

   In this case, $\boldsymbol{E}\cdot\hat{\boldsymbol{n}} = a$, which is indeed a constant. Bingo!

We've found that $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ is a constant in this case, so that's it. We're done! The integral becomes:

$$\Phi = \boldsymbol{E}\cdot\hat{\boldsymbol{n}}\int_{\substack{\text{curved surface}\\ \text{of cylinder}}} \mathrm{d}A$$

$$\Phi = a\int_{\substack{\text{curved surface}\\ \text{of cylinder}}} \mathrm{d}A$$

But the integral is just the curved surface area of the cylinder, which is $2\pi rL$, and so

$$\boxed{\text{Flux} = 2\pi raL}$$

> "*Consider a sphere of radius R in which there is a vector field of <u>changing</u> magnitude 1/r², which is everywhere pointing outwards from the centre of the sphere. What is the flux of the field through the surface of the sphere?*"

Diagram (to clarify things, I've only drawn one bit of the sphere, and one set of field lines):

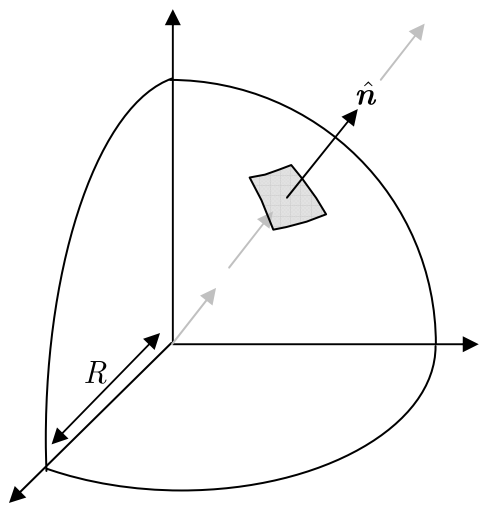

Let's go through the steps

1. **Decide how you're going to split up your surface**

   Clearly, in this case, it makes sense to break into small tiles, as shown above.

2. **Check if the flux through any bit of your surface is obviously 0**

   Nope ☹. There's stuff going through the whole thing…

3. **Find an expression for $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$**

   In this case, $\boldsymbol{E} = \frac{1}{r^2}\hat{\boldsymbol{r}}$ (because it's pointing radially outwards and has magnitude $1/r^2$). Furthermore, we saw that in the case of a sphere, $\hat{\boldsymbol{n}} = \hat{\boldsymbol{r}}$. Therefore, $\boldsymbol{E}\cdot\hat{\boldsymbol{n}} = 1/r^2$.

4. **Check that $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ isn't constant**

   Looking at the expression for $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$, you might be tempted to say that it's changing, because it includes the variable r. However, remember that we're only doing this integral over the surface of the sphere, where $r = R$ is constant! Therefore, $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ is constant.

So that's it! We're done! We can take $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ out of the integral, and get

$$\Phi = \boldsymbol{E}\cdot\hat{\boldsymbol{n}}\int_{\substack{\text{curved surface}\\ \text{of cylinder}}} \mathrm{d}A$$

$$\Phi = \frac{1}{R^2}\int_{\substack{\text{curved surface}\\ \text{of cylinder}}} \mathrm{d}A$$

The integral is, once again, just the surface area of a sphere, which is

$$\Phi = \frac{1}{R^2}4\pi R^2$$

$$\boxed{\text{Flux} = 4\pi}$$

OK, so far so good – we've managed to work out the flux from some fields. But in Gauss' Law, we have to do it the other way round – we know the flux, and we need the field. It's really, really, trivial to go from one to the other, but let's do it anyway just to get you completely comfortable with it… I'm going to take the case above, and just do it backwards

> "*A sphere of radius R lies in a constant vector field of unknown strength which is everywhere pointing outwards from the centre of the sphere. The strength of the field depends on r only. The flux of the field through that sphere is 4π. Find the unknown field*"

The diagram is exactly the same as above since we know the field is pointing radially outwards. This time, however, we don't actually know what the magnitude of the field is, so we'll just call it $E(r)$. Let's go through the seven steps.

1. **Decide how you're going to split up your surface**

   As above…

2. **Check if the flux through any bit of your surface is obviously 0**

   As above: Nope ☹. There's stuff going through the whole thing…

3. Find an expression for $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$

In this case, $\boldsymbol{E} = E\hat{\boldsymbol{r}}$ (because it's pointing radially outwards and has magnitude E). Furthermore, we saw that in the case of a sphere, $\hat{\boldsymbol{n}} = \hat{\boldsymbol{r}}$. Therefore, $\boldsymbol{E}\cdot\hat{\boldsymbol{n}} = E$.

4. Check that $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ isn't constant

Looking at the expression for $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$, you might be tempted to say that it's changing, because E is a function of r. However, remember that we're only doing this integral over the surface of the sphere, where $r = R$ is constant! Therefore, $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ is constant.

So we're done! $\boldsymbol{E}\cdot\hat{\boldsymbol{n}} = E(R)$ is constant, and we can just write

$$\Phi = \boldsymbol{E}\cdot\hat{\boldsymbol{n}}\int_{\substack{\text{curved surface}\\\text{of cylinder}}} \mathrm{d}A$$

$$\Phi = E(R)\int_{\substack{\text{curved surface}\\\text{of cylinder}}} \mathrm{d}A$$

$$\Phi = E(R)\times 4\pi R^2$$

However, we know that the flux is equal to 4π, because the question told us, and so

$$\Phi = E(R)\times 4\pi R^2 = 4\pi$$

$$E(R) = \frac{1}{R^2}$$

However, R is just a variable which we can take whatever value we want. Therefore, this is true for all radii, and

$$\boxed{\boldsymbol{E}(r) = \frac{1}{r^2}\hat{\boldsymbol{r}}}$$

The final example I'd like to try is one where you actually have to do the integral. I think it's beyond anything they might ask you in 8.02, but let's try it anyway, just to give you an idea of how they're done

> "*Consider a region of space in which there is a constant vector field* $\boldsymbol{E}(x,y,z) = a\hat{\boldsymbol{z}}$ *. A hemisphere of radius R lies on the x-y plane. What is the flux of the field through the sphere?*"

First, a diagram. Again, I've only drawn a quarter of the sphere and a single field line, for clarity:

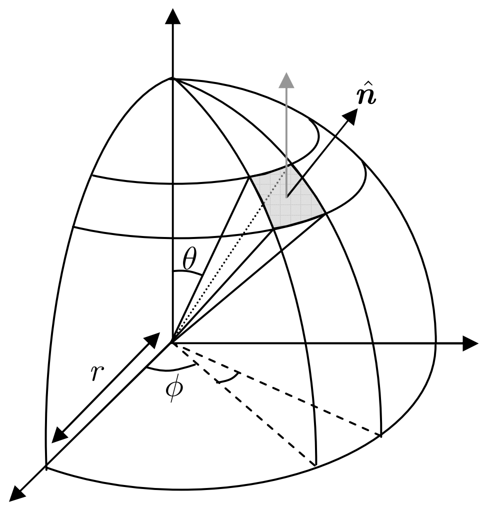

Let's go through the steps

1. **Decide how you're going to split up your surface**

   As usual, we split up the sphere into small tiles. (Note that it's only the hemisphere we want, so the "flat" bottom face is **not** included – if it was, what would you immediately be able to say the flux was?).

2. **Check if the flux through any bit of your surface is obviously 0**

   Nope – the field is everywhere on the curved surface ☹.

3. **Find an expression for** $\boldsymbol{E} \cdot \hat{\boldsymbol{n}}$

   To do this, let's draw a "tile" in greater more detail

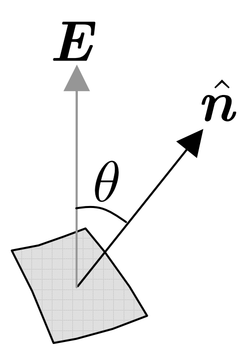

At this point, we're in trouble, because $\boldsymbol{E} = a\hat{\boldsymbol{z}}$ but $\hat{\boldsymbol{n}} = \hat{\boldsymbol{r}}$ (as it always does in spherical coordinates), but we don't know how to calculate the dot product $\hat{\boldsymbol{z}} \cdot \hat{\boldsymbol{r}}$, because they're in different coordinate systems.

The solution to this problem is to express $\hat{\boldsymbol{r}}$ in Cartesian coordinates. It turns out that $\hat{\boldsymbol{r}} = \sin\theta\cos\phi\hat{\boldsymbol{x}} + \sin\theta\sin\phi\hat{\boldsymbol{y}} + \cos\theta\hat{\boldsymbol{z}}$, as you might be able to see from this diagram (I'm afraid it's getting close to 4:00am, so the quality of my diagrams is deteriorating – apologies):

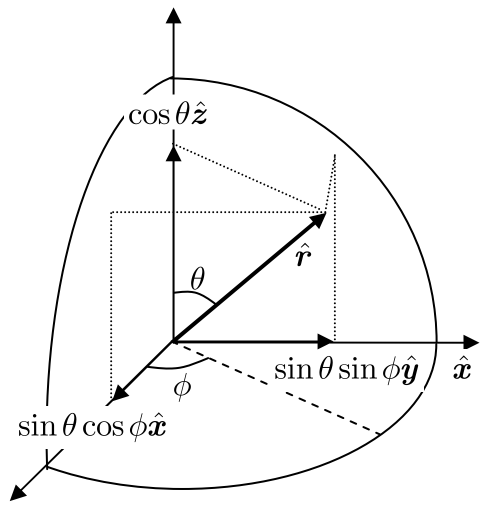

Therefore, we have that

$$\boldsymbol{E} \cdot \hat{\boldsymbol{n}} = a\hat{\boldsymbol{z}} \cdot \hat{\boldsymbol{r}} = a\hat{\boldsymbol{z}} \cdot \left(\sin\theta\cos\phi\hat{\boldsymbol{x}} + \sin\theta\sin\phi\hat{\boldsymbol{y}} + \cos\theta\hat{\boldsymbol{z}}\right) = a\cos\theta$$

4. **Check that $\boldsymbol{E} \cdot \hat{\boldsymbol{n}}$ isn't constant**

   No such luck this time! θ changes as we move to different places on the surface, so $\boldsymbol{E} \cdot \hat{\boldsymbol{n}}$ will also change. We need to integrate.

5. **Find dA for each little bit your surface**

   When we talked about tiling a sphere, we mentioned that each small time had area

$$\mathrm{d}A = R^2 \sin\theta \,\mathrm{d}\theta \,\mathrm{d}\phi$$

   Where R is the radius of the hemisphere.

6. **Set the correct limits on your integral**

So our integral at the moment is

$$\Phi = \int_{\substack{\text{surface of}\\\text{hemisphere}}} \boldsymbol{E}\cdot\hat{\boldsymbol{n}}\,\mathrm{d}A$$

$$\Phi = \int_{\substack{\text{surface of}\\\text{hemisphere}}} a\cos\theta\, R^2 \sin\theta\,\mathrm{d}\theta\,\mathrm{d}\phi$$

The only problem now is the phrase "surface of hemisphere" which we have instead of limits – that's not very mathematical! To convert that to math, we realise that there are two variables in this case - θ and ϕ. Let's see what the range of each variable has to be to encompass the sphere

- θ has to range from 0 to $\pi/2$ (from the z axis to the x-y plane)
- ϕ has to range from 0 to 2π (all around the z axis).

Therefore, the appropriate limits are

$$\Phi = \int_{\theta=0}^{\pi/2}\int_{\phi=0}^{2\pi} a\cos\theta\, R^2\sin\theta\,\mathrm{d}\phi\,\mathrm{d}\theta$$

[Note how I've swapped the order of ϕ and θ - that's just because there's no ϕ in the integral, so I can easily get rid of that integral first. Not essential, but it makes my life easier].

7. Solve it!

Time to get our integral solving skates out. First, take out the constants

$$\Phi = aR^2\int_{\theta=0}^{\pi/2}\int_{\phi=0}^{2\pi}\cos\theta\sin\theta\,\mathrm{d}\phi\,\mathrm{d}\theta$$

Then, perform the easy ϕ integration

$$\Phi = 2\pi aR^2\int_{\theta=0}^{\pi/2}\cos\theta\sin\theta\,\mathrm{d}\theta$$

Then, make the substitution $u=\sin\theta \Rightarrow \mathrm{d}u = \cos\theta\,\mathrm{d}\theta$

$$\Phi = 2\pi aR^2\int_0^1 u\,\mathrm{d}u$$

Do the easy integral

$$\Phi = 2\pi aR^2\left[\tfrac{1}{2}u^2\right]_0^1$$

And get

$$\boxed{\text{Flux} = \pi aR^2}$$

In retrospect, this is very re-assuring, because the flux through the bottom face of the hemisphere would have been $-\pi a R^2$ (Why? And why is there a minus sign?), which means that the two would sum 0. This makes sense, because the total flux through a closed surface (like the closed hemisphere) should be 0.

9. Gauss' Law

OK, enough with the math! Time for some physics! The whole point of going through this rigmarole was to be able to evaluate the surface integral in Gauss' Law:

$$\boxed{\oint_{\substack{\text{any closed}\\ \text{surface}}} \boldsymbol{E}\cdot\hat{\boldsymbol{n}}\,\mathrm{d}A = \frac{Q_{\text{enclosed}}}{\varepsilon_0}}$$

Gauss' Law is really incredibly easy to use – all you need to do is deal with the left hand side, then deal with the right hand side, and set them equal to each other ☺. What I want to do now is spend some time going through how to deal with each side, in detail…

10. The left-hand-side of Gauss' Law

You might think we've already done all the work needed to evaluated the LHS, because we know how evaluate surface integrals. This is nearly true, but there's a slight additional subtlety. Gauss' Law holds for <u>any</u> closed surface, so we're going to have to choose one particular surface over which to do an integral. In this section, I'm going to explain how to choose that surface.

There are only three steps to evaluating the LHS of Gauss' Law:

1. Determine what the field is going to look like (perhaps using symmetry).
2. Choose the Gaussian Surface that will make the surface integral as easy as possible (remembering the two conditions that can make a surface integral easy: "the flux through one of the faces might be 0" and "$\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ might be constant") This can only be one of three possibilities (all of which we've looked at in the last section)
    a. A box

b. A cylinder

c. A sphere

3. Perform the surface integral (in this case, $\boldsymbol{E}$ is still unknown, so you have to use the techniques in the second-to-last example above).

These are best shown by examples, in two sections' time…

11. The right-hand-side of Gauss' Law

The right-hand side of Gauss' Law is even easier – it really is as easy as finding the total charge enclosed inside the surface and dividing by ε_0.

The only subtlety involved is working out how much charge is contained in a **small part** of a shape. For example…

> "*A sphere of radius R has a charge Q uniformly distributed through it. How much charge is contained in the part of this bigger sphere with* $r < R$

The key to working this out is really very simple:

1. **Work out the charge density**

   In this case, the total charge is Q. The total volume is $\frac{4}{3}\pi R^3$. The total charge density is therefore $\rho = Q / \left(\frac{4}{3}\pi R^3\right) = 3Q / 4\pi R^3$

2. **Work out the volume of the smaller shape we're interested in**

   In this case, the smaller sphere has volume $\frac{4}{3}\pi r^3$.

3. **Multiply the two**

   Therefore, the total charge contained in the smaller shape is

$$q = \rho \times \tfrac{4}{3}\pi r^3$$

$$q = \frac{3Q}{4\pi R^3} \times \frac{4\pi r^3}{3}$$

$$\boxed{q = Q\frac{r^3}{R^3}}$$

The only reason this works is because the charge distribution is **uniform** throughout the sphere. It's possible for ρ to be different in different parts (in which case you'd need an integral), but I think this is slightly beyond the level of 8.02 (maybe not, though…)

12. Examples, examples, examples, examples, examples…

Phew! We're done! All that's left is applying it! Piece of cake ☺.

"*A sphere of radius R has a charge Q uniformly distributed through it. Find the field inside and outside the sphere*"

This is really two problems in one – first the inside of the sphere, then the outside. Let's do each separately

Outside the sphere

Let's go through our steps. First, the LHS of Gauss' Law

1. **Determine what the field is going to look like**

   What's the field going to look like outside a sphere? Well, we can use a symmetry argument to show that it must be of the form $\boldsymbol{E} = E(r)\hat{\boldsymbol{r}}$ [in other words, it must be radially outwards from the centre of the sphere, and have a value that depends only on its distance from the centre of the sphere]. The argument is as follows:

   - Let's imagine that $\boldsymbol{E}$ also had components in the $\hat{\boldsymbol{\theta}}$ and/or $\hat{\boldsymbol{\phi}}$ directions. What that would mean is that an electron resting on the surface of the sphere would be pushed left or right by the field. But that would imply that there's something on one side of the sphere pulling us that way. But that's impossible, because both sides of the sphere are identical, and anything pulling us on one side must also be pulling us from the other side, and cancelling out.

- Let's imagine that $\boldsymbol{E}$ also depended on θ and ϕ. That would mean that if we stood at a given point on the sphere, we'd feel something different if me moved left or if we moved right. But that's absolutely impossible, because both sides of the sphere are exactly the same.

Therefore, the field from the sphere is radially outwards:

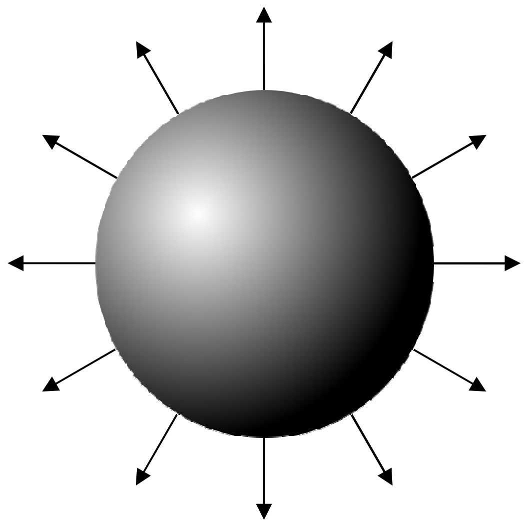

2. **Choose the Gaussian Surface that will make the surface integral as easy as possible**

Let's look at the two possible simplifications to the surface integral and see if we can choose a Gaussian surface that fulfils either one of them:

- **Making the flux through one of the faces 0** – can't think of any easy way to make this happen…
- **Making $\boldsymbol{E}\cdot\hat{\boldsymbol{n}}$ constant** – this looks promising… The field lines are all radially outwards, and the field is **constant** at a given radius. So it seems a good Gaussian surface will (1) have $\hat{\boldsymbol{n}}$ also pointing radially outwards (2) all be located at a single radius.

A surface that satisfies such conditions is the surface of a **sphere** – it's always at a single radius, and $\hat{\boldsymbol{n}}$ points radially outwards.

How big should the sphere be, though?

A very common mistake is to assume that the Gaussian surface needs to be physical (so, for example in this case, that the Gaussian surface needs to be the actual surface of our physical sphere, and that the radius of our Gaussian surface therefore needs to be R). However, this is *not true*. The Gaussian surface is a **fictitious** mathematical object, which we just create to use Gauss' Law.

So, how big does it have to be? Well, it just depends where we want to find $\boldsymbol{E}$. If we want to know what $\boldsymbol{E}$ is at a distance X from the centre of the physical sphere, then we need to make sure the field lines *at that distance* cross our Gaussian surface. To do that, we give our Gaussian surface a radius X.

In this case, we're only interested in the field *outside* the physical sphere, so we'll simply say our Gaussian surface has radius r, where $r > R$:

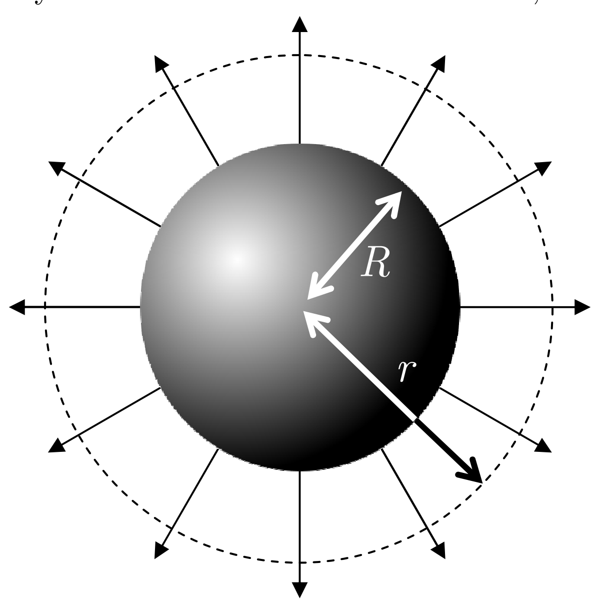

3. Perform the surface integral

Finally, we need to perform the integral for the field $\boldsymbol{E} = E(r)\hat{\boldsymbol{r}}$ over the sphere of radius r. Conveniently, we did this in the previous section, and came up with

$$\int_{\substack{\text{surface of sphere}\\\text{of radius } R}} \boldsymbol{E}\cdot\hat{\boldsymbol{n}}\,\mathrm{d}A = 4\pi r^2 E(r)$$

Now let's deal with the RHS of Gauss' Law. In this case, it's simple, because whatever r is, it'll always include the entire sphere (because $r > R$, since we're looking for the field *outside* the sphere) and so the enclosed charge is just Q. Therefore, the RHS is just Q / ε_0.

Putting both sides together, Gauss' Law gives

$$E(r)\times 4\pi r^2 = \frac{Q}{\varepsilon_0}$$

$$\boxed{E(r) = \frac{Q}{4\pi\varepsilon_0 r^2}}$$

This holds for *any* r outside the sphere, and so we've solved the first part of the problem!

Inside the sphere

Inside the sphere, absolutely *everything* we've said above still applies, except for two things:

- Our Gaussian surface will now need to be chosen for $r < R$, since we're interested by stuff *inside* the sphere:

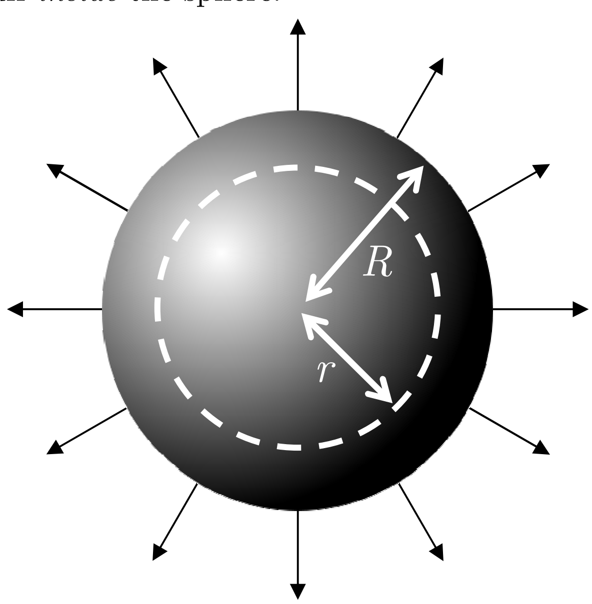

- The integral for flux will **not** change in form, because only field lines from stuff **inside** the Gaussian surface will matter. The field lines from outside will just enter the sphere and leave again, adding 0 to the flux.
- The charge contained inside the Gaussian surface will now be different. Assuming our Gaussian surface has a radius r, we can use the result of the previous section to find that the charge contained in it is
$$q = Q\frac{r^3}{R^3}$$
And so the RHS of Gauss' Law will simply be $Qr^3 / \varepsilon_0 R^3$

Applying Gauss' Law to everything we said above, we get

$$E(r) \times 4\pi r^2 = \frac{Qr^3}{\varepsilon_0 R^3}$$
$$\boxed{E(r) = \frac{Qr}{4\pi\varepsilon_0 R^3}}$$

And this is the solution to the second part of the problem!